

-1-

OBJECT-ORIENTED PROCESSOR DESIGN AND DESIGN
METHODOLOGIES

BACKGROUND OF THE INVENTION

5 This invention relates generally to methods for object-oriented hardware design and to hardware produced thereby, and more particularly to distributed memory, object-oriented, class-based methods for processor design and the processors produced thereby.

10 Voice over Internet protocol (VoIP), wideband code division multiple access (WCDMA), third generation wireless networks, and other advanced wireless and wired broad-band communication systems require many complex, computationally-intensive signal processing functions. Examples of such functions include orthogonal frequency division multiplexing (OFDM) modems, Viterbi decoders, and Reed-Solomon decoders. In many cases, it is desirable for these signal processing functions to be fabricated on a single-chip integrated circuit.

15 One known methodology for placing such highly complex systems on a single chip is to provide powerful computational platforms on the chip to process all functions in a sequential manner in conjunction with a number of tightly-coupled intellectual property (IPs) cores, i.e., special-purpose processor and firmware layouts that are licensed for use in chip layouts for more complex processors.

20 Computational platforms used in this design methodology include one or more microprocessors or digital signal processors (DSPs) and one or more standard or proprietary communication backbones, interface buses, or virtual sockets to connect all the necessary components into a unified environment. Such platforms can be characterized as being "processor-centric," because the various processors and IP

25 cores share complex bus architectures to process all of the functions and algorithms sequentially. As systems become larger and more complex, even more powerful core processors are required.

Although presently known signal processing architectures and design methodologies are sufficient for many present applications, it is becoming

-2-

increasingly difficult to meet processing demands of new applications with these architectures for several reasons. First, newer bus and controller architectures have become very complicated because memory speed cannot keep up with the increasing speed of central processing units (CPU), even with cache memory. Thus, there is a CPU-memory bottleneck that manifests itself in faster applications due to physical propagation factors whenever the silicon die area used by a processor is sufficiently large. Second, present architectures require a costly investment in very large and complex software. When suitable software is written, it is necessarily operating system (OS) dependent, because such dependency is required to ensure that each process receives an appropriate time slice of the CPU's computational resources. Whenever changes to a processor are necessary or a move is made to another OS platform, the prior effort and investment in developing the application software are largely wasted or rendered obsolete. Third, computational demands on processor-centric architectures require increased computational speed as the processes themselves become more complex. Increases in computational speed necessarily raise power consumption.

It would therefore be desirable to provide methods for designing complex application processors that avoid CPU memory bottlenecks due to large silicon die areas. It would also be desirable to provide a processor architecture that provides reduced dependence upon an operating system of any particular core processor and a design method that provides greater freedom to redesign the application processor around a different host processor. It would also be desirable to provide an application processor architecture in which smaller functional sets not requiring a single high-speed processor are performed relatively independently of one another, thereby avoiding the CPU-memory bottleneck.

BRIEF SUMMARY OF THE INVENTION

There is therefore provided, in one embodiment of the present invention, a distributed processing system having a host processor including a host

-3-

communication infrastructure (HCI) configured for communication with said host processor; a plurality of class processors each having an associated private localized read/write memory; and a plurality of application program interface modules each configured to provide an interface between said host communication infrastructure
5 and at least one said class processor, wherein each said class processor responds to selected data messages on said HCI to perform selected computations utilizing said read/write memory. This embodiment provides an ideal architecture for fabrication on a single chip. This embodiment also avoids processor and bus bottlenecks by providing distributed processing power with local memory for each class processor.

10 There is also provided, in another embodiment of the present invention, a method for designing a distributed processing system for an application. The method includes steps of partitioning the application into functions and data messages; configuring a host processor having a host communication infrastructure (HCI) to pass data messages via the HCI to control the application; configuring a
15 plurality of class processors to compute the functions into which the application is partitioned in response to the data messages; and interconnecting the class processors to the host processor via application program interface modules in a star configuration. Systems designed in accordance with this method embodiment are well-suited for integration on a single chip, and can be easily updated and modified
20 as necessary, because changes made to a class processor have minimal effect on the remainder of the system.

Other advantages of these embodiments and the others disclosed herein will become apparent to those skilled in the art upon reading the detailed description in conjunction with the accompanying drawings.

25 BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 is a drawing of a block diagram of an embodiment of a class processor of the present invention.

Figure 2 is a drawing of a block diagram of an embodiment of a physical layer processor of the present invention.

Figure 3 is a drawings of a block diagram of an embodiment of an application layer processor of the present invention.

5 DETAILED DESCRIPTION OF THE INVENTION

As used herein, the term "object-oriented" refers to a paradigm in which variables and command statements operating on the variables are termed "objects," variables associated with an object are called "attributes," and functions that operate on them are termed "methods." "Encapsulation" refers to a feature provided by the object oriented paradigm in that the only way to operate on, view, or otherwise
10 alter, read, or access attributes of an object is by invoking the object's methods.

In one embodiment of the present invention, distributed processing system hardware is designed in accordance with the object-oriented paradigm. An application, for example, a physical layer processor for a communication system, is
15 divided into active entities, i.e., signals, processing units, and defined transformations of signals. A correspondence is drawn between signals and objects and between transformations and methods, with transformations being defined as mathematical operations that are performed on signal objects. In cases in which it is possible to divide the application into active entities in more than one manner, a
20 division is selected that localizes resources needed for implementation so that communication with other functions is reduced or minimized. This criterion reduces a potential for system conflicts and also reduces timing overhead for coordinating shared resources.

For example, a physical layer processor of a type suitable for some
25 applications is partitioned according to this method embodiment into (1) an A/D converter for converting an analog signal into a time domain digital signal "A," (2) a standalone FFT using radix-2 to transform time domain digital signal "A" into a frequency domain signal "B," the standalone FFT processor including twiddle factor

-5-

coefficients and many repeated operations; (3) a Viterbi or convolutional decoder to decode signal "B" into a signal "C" using a read/write memory; (4) a standalone FFT using a radix-4 transform to convert signal "C" to signal "D" and including other twiddle factor coefficients and many repeated operations; and (5) a multiplexer to take signal "D" to a host processor in the form of another digital signal "E." It follows from the above partitioning and from the definition of an object, that signals A, B, C, D, and E are objects. It also follows that an analog to digital conversion, a fast Fourier transformation, a convolutional decoding, and a multiplexer mapping operation are methods. Computational engines that implement the methods are also identified as objects. Twiddle factors and memory are identified as attributes of objects, because they are attributes of the computational engines.

Groups of related functions and objects are then selected with an objective of minimizing communication between functional groupings, although it will be understood that not all application processors will be broken down into functions and objects that can be grouped in this manner. For example, a Fourier transform class providing an FFT and a DFT is recognized as a class grouping. Each computational object is configured to respond to requests from at least a host processor, such requests being defined as predefined or selected data messages that differ for each function. The transfer of messages between function processors, which henceforth are referred to as "class processors", allows relatively intense computation to be performed by the class processors without excessive loading of a host processor data bus.

Referring now to Figure 1, an exemplary embodiment of a class processor of the present invention comprises a special purpose processor 12. Suitable special purpose processors 12 include, but are not limited to, a very short instruction word (VSIW) processor, a digital signal processor (DSP), an application specific integrated circuit (ASIC) processor, a suitably-programmed microprocessor subcomponent, and hardwired components. In one embodiment in which class processor 10 is an FFT processor, special purpose processor 12 is a VSIW DSP

processor. In one embodiment in which class processor 10 is a RAKE finger receiver, special purpose processor 12 is an ASIC processor. In one embodiment, special purpose processor 12 includes software or firmware (not shown) to provide a portion of its functionality.

5 Depending upon functions to be performed by class processor 10, one or more localized read/write memories 14, 16, and 18 are provided and interconnected with special purpose processor 12 so as to be directly accessible to it, i.e., visible and addressable in its memory space. Memory 14 is a private localized read/write memory 14 that is used by and is accessible only by special purpose processor 12
10 through its implemented methods. Memory 16 is a localized protected read/write memory that is used by special purpose processor 12 to store and/or read data accessible only by other class processors 10 of the same grouping or class. For example, where class processor 10 is a fast Fourier transform (FFT) processor, another class processor 10 (not shown in Figure 1) that implements a similar
15 function or functions is able to access protected memory 16. Examples of two class processors sharing protected localized read/write memory 16 are two FFT processors used at the same time, or an FFT processor and a discrete Fourier transform (DFT) processor that computes discrete Fourier transforms somewhat differently, for example, using a Winograd DFT. Access to protected localized
20 read/write memory 16 is provided by one or both of a direct interconnection 18 from protected localized read/write memory 16 to the other class processor 10 or by providing special purpose processors 12 of each class processor 10 belonging to the same class with special knowledge of messages that can be passed between class processors 10 of that class. Memory 20 is a public read/write memory that can be
25 addressed by other components, including a host processor (not shown in Figure 1). In one embodiment, public read/write memory is directly accessible via a memory mapped mailbox or an application program interface (API) module 22. It will be recognized by those skilled in the art that not all embodiments of class processors 10 require all three types of memories 14, 16, and 20. API module 22 comprises a

-7-

communication port, for example, a hardware or software communication port that including a memory-mapped dual port memory bank. In another embodiment, API module 22 includes a memory stack.

In another embodiment and referring to Figure 2, a block diagram of an exemplary physical layer communications processor 24 is shown. In this example, processor 24 is an object oriented communications signal processor (OOCSP) fabricated on a single chip. OOCSP 24 comprises a host processor 26 and various functional class processors 26, 28, 30, 32, and 34 in addition to an embodiment of FFT class processor 10 of Figure 1. Host or main processor 24 is a standard microprocessor embodiment or "IP core." The invention places no restriction on the type of host processor 24. Exemplary host processors 24 useful for this embodiment are ARM, MIPS, x86, 68xxx, TMS320, DSP16xxx and DSP56xxx series processors. Each of these host processors is characterized by a native host communication infrastructure (HCI) 36, which includes a bus and a port configuration.

Class processors 10, 26, 28, 30, 32, and 34 provide computational power for invoking methods of data or signal objects for OOCSP 24. For the embodiment exemplified by Figure 2, the class processors are implemented as hard IPs and are FFT class processor 10, Galois field (GF) field class processor 26, synchronization object 28, forward error correction (FEC) object 30, Rake finger object 32, and I/O class processor 34, the latter being provided to increase data pumping capabilities of OOCSP 24 beyond that already provided by native I/O processing of host processor 26. Class processors 10, 26, 28, 30, 32, and 34 are each designed to processes certain specialized functions and no others on data structures, or in other words, each performs operations on a selected proper subset of application objects. Each class processor 10, 26, 28, 30, 32, and 34 is an application oriented functional unit programmable for its intended application. For example, FFT class processor 10 can initiate various real and complex FFT operations at different word-length accuracy, and is the hardware analog to an object-oriented language class. The

-8-

functions of class processors 10, 26, 28, 30, 32, and 34 are selected and grouped to enhance abstraction, i.e., making the action of OOCSP 24 more readily accessible to a programmer of host processor 26, and to enhance encapsulation, i.e., hiding of internal workings of each class. Memory 38 is also provided for programming and local data storage of host processor 26. For efficiency and speed, memory 38 is organized in a native structure of host processor 26, whether it is provided internally or externally to host processor 26. Memory 38 is of a suitable type (e.g., RAM, ROM, DRAM, eDRAM, etc.) and amount needed for host processor 26 to control the functions of class processors 10, 26, 28, 30, 32, 34 and their communication with host processor 26.

In the embodiment of Figure 2, APIs 22 provide communication between class processors 10, 26, 28, 30, 32, and 34 and other subsystems, including host processor 26 and other class processors. Each API 22 provides "public visibility" by providing an interface to HCI 36. Each API 22 provided to a different class processor 10, 26, 28, 30, 32, and 34 implementing different functions is slightly different, in that each provides a visibility to host processor 26 for its respective class processor that effectively defines a programming interface for its class processor. Thus, the API of a class effectively describes what a class can do, while implementation of the class describes how it does it, in a manner analogous to application programming interfaces of software programs. Particular implementation of APIs 22 is a design choice, in that APIs 22 can be a virtual socket interface or any type of bus communication or I/O port data exchange mechanism. Use of a memory mapped dual port RAM bank, stack, and FIFO memory permits native HCI 36 to be easily maintained.

It will be observed that OOCSP 24 provides a number of advantages over known processors. First, communication among class processors 10, 26, 28, 30, 32, and 34 is restricted, in that in the embodiment of Figure 2, no class processor has any knowledge of any other class processor or any ability to communicate with any other class processor. All communication is with host processor 26, i.e., OOCSP

24 is arranged in a "star" configuration. No complex, fast busses are needed. Because class processors 10, 26, 28, 30, 32, and 34 each provide separately implemented functions with a defined interface, no class processor has any particularized knowledge of the internal workings of any other class processor, and
5 class processors are referred to only through their defined interfaces. Each class processor has all the memory it requires in close proximity to itself to reduce propagation delays, and the minimization of public or global memories and data structures is kept to a minimum to limit the opportunity for class processors to affect one another. Because of abstraction and data hiding behind a defined interface, class
10 processors 10, 26, 28, 30, 32, and 34 can be provided as hard IP objects that can be used without concern as to their inner workings. Moreover, any changes to one of the class processors has only minimal, if any, effect on the operation of others, so the effect of design revisions is localized. Also, class processors that are closely coupled to the host processor or to other class processors (i.e., that require frequent
15 or speedy, low propagation delay access to one another) can be placed close to one another on a chip.

In another embodiment and referring to Figure 3, a block diagram of a single-chip OOCSP application layer processor 40 is shown. Although designed for a different purpose, the architecture of application layer processor 40 is similar to
20 that of physical layer processor 24 of Figure 2, except that host processor 26 and HCI 36 communicate with class processors 42, 44, 46, 48, 50, and 34, most of which are different from those comprising physical layer processor 24. More particularly, application layer processor 40 comprises an audio class processor 42, a filter bank class processor 44, a synchronization object 46 a video class processor
25 48, a discrete cosine transform (DCT) object and an I/O class processor 34. This example shows the reuse of I/O class processor 34, which is made possible, in part, because of the loose coupling, i.e., data hiding and encapsulation, of the object-oriented design paradigm. Synchronization object 46 is related to synchronization object 28, but differ because synchronization methods applicable to a physical layer

and to an application layer are, in general, somewhat different. Class processors 42, 44, 46, 48, 50 and 34 are arranged so that those functions that would be most adversely affected by propagation delays are electrically closest to host processor 26 in that propagation delays are minimized. For further efficiency, although not
5 shown in the embodiments of Figures 2 or 3, related class processors, i.e., those of the same class, share a protected localized read/write memory 16 in other processor embodiments. Sharing is accomplished in one embodiment by direct coupling of memory 16 via a semi-private bus 18 to the sharing class processor or processor, or in another embodiment by configuring processors 12 of class processors of the same
10 class to communicate messages to one another through their APIs 22, either by direct communication with one another or indirectly through a host processor 26. Otherwise, class processors have no knowledge of communication protocols of class processors not in the same class and thus, cannot communicate with or reference the other processors, and memories 16 are not directly or indirectly accessible to class
15 processors of different classes. In one embodiments, class processors are provided with public memories 20 that are addressable by host processor 26 via HCI 36 and API 22, such as by memory mapping.

From the preceding description of various embodiments of the present invention, it is evident that complex processors are produced from the design
20 methodologies of embodiments of the present invention without CPU memory bottlenecks due to large silicon die areas, because of data hiding and the providing of localized private memories for class processors. Moreover, the application processor architecture resulting from embodiments of the present invention isolate host processors from the complexity of the functions provided by the class
25 processors with an application programming interface, so that the resulting architectures are relatively independent of an operating system running on the host processor. Also, the functional decomposition of the application processor allows greater freedom to change host processors and removes high speed processing

constraints by shifting much of the processing load to specialized processors rather than the host processor.

Although the invention has been described and illustrated in detail by reference to specific embodiments, it is to be clearly understood that the same is intended by way of illustration and example only and is not to be taken by way of limitation. For example, one skilled in the art will recognize that one can build up a collection of reusable objects, classes, and class hierarchies for different applications, their reuse being facilitated by their appearance to a programmer of a host processor as simple API module calls. Reuse and design validation is further enhanced and facilitated because software simulators of object-oriented classes and their associated objects can be made identical in behavior to the processors themselves to allow rapid and accurate software development and final compilation and integration of an overall device. Indeed, devices designed using method embodiments of the present invention can be implemented using technologies analogous to software compilers. It will also be observed that the invention is generally applicable to many different applications, including, for example, user-layer applications such as object-oriented HDTV processors and object-oriented audio processors, and also to applications not related to signal processing, as such. Because of the modularity of design provided by the present invention, the present invention can be used in the design of, and incorporated into the architecture of "super" OOCSP platforms in which different OOCSP cores are combined. Accordingly, the spirit and scope of the invention are to be limited only by the terms of the appended claims.